

CLAIMS

1. A method of restricting propagation of viruses in a network having a plurality of hosts, comprising the steps of:

monitoring network activity of a first host of the plurality and establishing a first record which is at least indicative of identities of hosts within the network contacted by a first host;

limiting contact of the first host to other hosts within the network over the course of a first time interval, so that during the first time interval the first host is unable to contact more than a predetermined number of hosts not in the first record;

wherein the method further comprises an additional selection process for determining which hosts of the plurality the first host is allowed to contact.

2. A method as claimed in claim 1, wherein the first record is indicative of identities of hosts within the network either: to whom data has been sent by the first host; or with whom a socket has been established.
3. A method as claimed in claim 1, further comprising the step of establishing a second record which identifies destination hosts with whom contact is a priority; and

wherein during the first time interval the first host is unable to send data to more than a predetermined number of hosts not in both the first record and the second record.

4. A method as claimed in claim 3, wherein the second record is indicative of connection characteristics of priority requests; and

wherein during the first time interval the first host is unable to send data to more than a predetermined number of destination hosts which are not in the first record and which do not match priority request characteristics and the host identities of the second record.

5. A method as claimed in claim 3, wherein the second record is established by a user.

6. A method as claimed in claim 3, wherein the second record is established by examining the system configuration of the first host.

7. A method as claimed in claim 1, further comprising the steps of:

diverting requests to contact hosts not in the first record to a delay buffer; and

transmitting the predetermined number of requests from the delay buffer at the end of the first time interval.

8. A method as claimed in claim 7, further comprising the step of determining a request characteristic indicative of at least one of the origin or the protocol of each request within the delay buffer.

9. A method as claimed in claim 8, further comprising the step of selecting the predetermined number of requests for transmission from the delay buffer by:

calculating a number of requests sharing each determined request characteristic;

determining the request characteristic shared by the lowest number of requests;
and

selecting for transmission requests from the delay buffer those requests which have common the least common request characteristic.

10. A method as claimed in claim 9, further comprising the steps of:

determining the next least common request characteristic; and

selecting the requests sharing the next least common characteristic.
11. A method as claimed in claim 8, further comprising the steps of:

calculating a number of requests sharing each request characteristic; and

checking if the number of requests sharing a characteristic rises above a predetermined threshold, then blocking all requests sharing that characteristic.
12. A method as claimed in claim 11, wherein the threshold corresponds to 50% of the total number of requests in the buffer if the total number of requests exceeds a predetermined minimum threshold.
13. A method as claimed in claim 11, wherein in the check on the number of requests, a separate predetermined threshold is applied for each respective request characteristic.
14. A method as claimed in claim 8, further comprising the steps of:

calculating the number of requests sharing each request characteristic; and

if the number of requests sharing a single characteristic rises above a predetermined threshold, then removing from the delay buffer all requests sharing that characteristic.

15. A method of operating a host in a network comprising the steps of:
monitoring requests to send data to other ("destination") hosts in the network over the course of successive time intervals;

delaying transmission of requests which are outside the scope of a policy;

establishing a buffer for storing data relating to delayed requests, the requests being stored in the buffer in correspondance with time of receipt; and

selecting from the buffer for transmission requests conforming to a predetermined criterion other than the order in which the requests are stored in the buffer.
16. A method according to claim 15 wherein the buffer stores a plurality of entries each, entry identifying a request.
17. A method according to claim 16 wherein the requests per se are diverted to the buffer, and form the buffer entries.
18. A method according to claim 16 wherein the buffer entries indicate a request characteristic identifying request origin, or request protocol.
19. A method according to claim 18 wherein the predetermined criterion selects for transmission requests which share the least common request characteristic.
20. A method according to claim 15 wherein the predetermined criterion selects requests identifying destination hosts on a priority list.
21. A memory storing a computer program for providing instructions for a first host in a network to monitor its activity and to establish a record which is at least indicative of identities of other hosts within the network contacted by the first host, and to limit the ability of the first host to contact other hosts within the

network over the course of a first time interval, so that during the first time interval the first host is unable to contact more than a predetermined number of other hosts not in the first record, wherein the program further includes instructions for instructing the first host to perform an additional selection process for determining which hosts of the plurality the first host is allowed to contact.

22. A computing entity in a network, programmed to establish, by monitoring its activity, a record indicating identities of other entities in the network to whom a communication has been sent, and, for the duration of a predetermined interval of time, to restrict dispatch of communications to other network entities whose identity is not in the record, the computing entity further being programmed to perform an additional selection process to determine to which network hosts whose identity is not in the record it is allowed to dispatch a communication.
23. An entity according to claim 22 wherein the additional selection process selects a predetermined number of requests to dispatch a communication to an entity not identified in the record.
24. An entity according to claim 23 wherein the additional selection process operates to: calculate how many requests to dispatch a communication share a particular characteristic; determining the request characteristic shared by the lowest number of requests ("least common request characteristic"); and select for transmission those requests which have in common the least common request characteristic.
25. An entity according to claim 22 programmed to store requests to dispatch communications to network entities whose identity is not in the record in a buffer.
26. An entity according to claim 25 wherein the entity is programmed to perform the additional selection process on requests stored in the buffer.

27. An entity according to claim 26 wherein requests are stored in the buffer in temporal order, and the additional selection process operates to re-order requests stored in the buffer.
28. An entity according to claim 27 wherein a predetermined number of requests to dispatch communication to entities whose identity is not in the record are dispatched from the buffer with the passing of each predetermined interval of time.
29. A network having a plurality of computing entities according to claim 22.